

WHAT IS CLAIMED IS:

1 1. In a disk-based data storage system, a method for reference
2 count regeneration, the method comprising:

3 sorting a virtual track table (VTT) to generate a reference list
4 wherein the reference list contains each reference made by an entry in the VTT to
5 an entry in a track number table (TNT);

6 counting the number of references in the reference list that refer to
7 each unique entry in the TNT to generate a respective reference count list; and
8 merging the reference count list with the TNT to generate an updated
9 TNT.

1 2. The method of claim 1 further comprising allocating a first
2 memory region to hold the reference list, and allocating an additional memory
3 region to hold further references that comprise the respective reference list when
4 the sorting encounters an end of the first memory region.

1 3. The method of claim 1 wherein the sorting is a radix sorting
2 that is performed using a bitwidth that is selected based on at least one disk cache
3 attribute.

1 4. In a disk-based data storage system, a method for reference
2 count regeneration, the method comprising:

3 radix sorting a virtual track table (VTT) to generate a reference list
4 that comprises a plurality of sub-lists wherein each of the sub-lists contains
5 references to a unique range of entries in a track number table (TNT);

6 counting the references in the reference list that refer to each unique
7 entry in the TNT to generate a respective reference count list; and

8 merging the reference count list with the TNT to generate an updated
9 TNT.

1 5. The method of claim 4 wherein the radix sorting comprises:
2 reading the VTT;

3 writing a first set of the sub-lists wherein each of the first set of
4 sub-lists contains references to the TNT entries within a member of a first set of the
5 unique ranges of entries in the TNT;

6 reading a selected one of the first set of sub-lists; and
7 writing a subsequent set of sub-lists wherein each member of the
8 subsequent set of sub-lists contains references to entries within a member of a
9 subsequent set of ranges of entries in the TNT, and the subsequent set of ranges
10 comprises ranges that are within the range of the first set sub-list being read, and
11 wherein the steps of reading the first set of sub-lists and writing the subsequent set
12 of sub-lists are performed in parallel over a non-empty group of the first set of
13 sub-lists.

1 6. The method of claim 4 wherein the radix sorting comprises
2 at least one first pass radix sort that comprises:

3 reading a first unique portion of the VTT;
4 writing a first set of sub-lists wherein each of the sub-lists contains
5 references to TNT entries within a member of a first set of the unique ranges of
6 entries in the TNT;
7 reading a subsequent unique portion of the VTT; and
8 writing a subsequent set of sub-lists wherein each of the sub-lists
9 contains references to TNT entries within a member of the first set of the unique
10 ranges of entries in the TNT, and wherein the steps of reading the first unique
11 portion of the VTT, reading the subsequent unique portion of the VTT, writing the
12 first set of sub-lists, and writing the subsequent set of sub-lists are performed in
13 parallel over elements of non-empty groups of the first and the subsequent unique
14 portions of the VTT.

1 7. The method of claim 6 further comprising at least one
2 subsequent pass radix sort that comprises:

3 reading the first set and the subsequent set of sub-lists; and
4 writing respective next subsequent sets of sub-lists wherein members
5 of each of the next subsequent sets of sub-lists contain references to TNT entries
6 within respective members of a subsequent set of ranges of entries in the TNT, and

7 wherein the steps of reading a sub-list and writing the next subsequent set of
8 sub-lists are performed in parallel over elements of the non-empty group of
9 sub-lists.

1 8. The method of claim 7 further comprising:
2 counting references in each respective next subsequent set of sub-lists
3 to generate a respective partial reference count list; and
4 merging the partial reference count lists to generate the reference
5 count list.

1 9. The method of claim 8 wherein the step of counting the
2 number of references in each respective next subsequent set of sub-lists is
3 performed substantially simultaneously over elements of a non-empty group of the
4 next subsequent set of sub-lists.

1 10. The method of claim 4 wherein the steps of counting and
2 merging are performed in parallel over elements of a non-empty group of the
3 sub-lists.

1 11. The method of claim 6 wherein a union of the members of the
2 first set of unique ranges of entries in the TNT contains all of the entries in the
3 TNT.

1 12. The method of claim 5 wherein the step of writing comprises:
2 allocating space in a memory region to hold a portion of a selected
3 one of the sub-lists; and
4 comparing the size of the selected sub-list to the size of the space
5 allocated to hold the selected sub-list, and allocating additional space to hold a
6 subsequent portion of the selected sub-list when the size of the selected sub-list is
7 equal to or greater than the size of the space allocated to hold the selected sub-list.

1 13. The method of claim 4 further comprising measuring a
2 reference count list length, the measuring comprising:

3 selecting a subsequent set of ranges of entries in the TNT to be
4 generated by a subsequent radix sort pass;

5 counting the number of references to TNT entries within each of the
6 members of the subsequent set of ranges; and

7 allocating space to hold a sub-list to be generated by the subsequent
8 radix sort pass, wherein the space is large enough to hold the number of references
9 to the TNT entries within the member of the subsequent set of ranges that
10 corresponds to the sub-list generated by the subsequent radix sort pass.

1 14. The method of claim 4 wherein the merging comprises:
2 reading a reference count value from the reference count list;
3 reading a TNT entry that corresponds to the reference count value;
4 comparing the reference count value contained within the TNT entry
5 to the reference count value from the reference count list; and

6 writing an entry into a reference count mis-compare list, replacing
7 the reference count value within the TNT entry with the reference count value from
8 the reference count list, and writing the updated TNT entry back into the TNT
9 when the reference count value contained within the TNT entry is not equal to the
10 reference count value from the reference count list.

1 15. The method of claim 14 wherein the merging further
2 comprises refraining from writing unchanged TNT entries back into the TNT when
3 the reference count values within a block of contiguous TNT entries are equal to
4 the corresponding reference count values from the reference count list.

1 16. The method of claim 15 further comprising refraining from
2 writing the reference count mis-compare entries into the reference count
3 mis-compare list when the number of entries in the reference count mis-compare
4 list is equal to or greater than a first predetermined number of incorrect reference
5 counts.

1 17. The method of claim 15 further comprising refraining from
2 writing the reference count mis-compare information into the reference count

3 mis-compare list when the time interval for writing the reference count mis-
4 compare values into the reference count mis-compare list exceeds a predetermined
5 time or the size of the mis-compare values exceeds a predetermined size in the
6 memory region.

1 18. In a disk-based data storage system, a system for reference
2 count regeneration, the reference count regeneration system comprising:

3 a sorter configured to sort a virtual track table (VTT) to generate a
4 reference list wherein the reference list contains each reference made by an entry
5 in the VTT to an entry in a track number table (TNT);

6 at least one counter configured to count the number of references in
7 the reference list that refer to each unique entry in the TNT to generate a respective
8 reference count list; and

9 a merger configured to merge the reference count list with the TNT
10 to generate an updated TNT.

1 19. The reference count regeneration system of claim 18 further
2 comprising a first memory region allocated to hold the reference list, and an
3 additional memory region allocated to hold further references that comprise the
4 respective reference list when the sorter encounters an end of the first memory
5 region.

1 20. The reference count regeneration system of claim 18 wherein
2 the sorter is a radix sorter that is configured to sort using a bitwidth that is selected
3 based on at least one disk cache attribute.